

TEACHING LLMs TO ABSTAIN VIA FINE-GRAINED SEMANTIC CONFIDENCE REWARD

Anonymous authors

Paper under double-blind review

ABSTRACT

Mitigating hallucinations in Large Language Models (LLMs) is critical for their reliable deployment. Existing methods typically fine-tune LLMs to abstain from answering questions beyond their knowledge scope. However, these methods often rely on coarse-grained signals to guide LLMs to abstain, such as overall confidence or uncertainty scores on multiple sampled answers, which may result in an imprecise awareness of the model’s own knowledge boundaries. To this end, we propose a novel reinforcement learning framework built on **Fine-grained Semantic Confidence Reward (FISCORE)**, which guides LLMs to abstain via sample-specific confidence. Specifically, our method operates by sampling multiple candidate answers and conducting semantic clustering, then training the LLM to retain answers within high-confidence clusters and discard those within low-confidence ones, thereby promoting accurate post-hoc abstention. Additionally, we propose a new metric for evaluating the reliability of abstention fine-tuning tasks more comprehensively. Our method significantly enhances reliability in both in-domain and out-of-distribution benchmarks.

1 INTRODUCTION

Large language models (LLMs) have achieved remarkable success, demonstrating powerful capabilities in content generation, complex reasoning, and software development (Achiam et al., 2023; Grattafiori et al., 2024; Yang et al., 2024). Their widespread adoption has established them as essential tools in numerous applications. However, a critical vulnerability threatens their reliability: their propensity to hallucinate, fabricating plausible but false information when faced with questions beyond their knowledge (Zhang et al., 2023). Such hallucinations erode user trust and can propagate misinformation. Reliance on fabricated content in high-stakes domains such as medicine or law can have severe consequences. Therefore, mitigating hallucinations is essential for the safe and responsible deployment of LLMs.

A promising direction for mitigating hallucination is to fine-tune LLMs to abstain from answering questions that lie beyond their knowledge scope and answer those within it (Wen et al., 2025; Cheng et al., 2024). Supervised methods often categorize a training set into “known” and “unknown” questions based on the correctness or accuracy of the model’s answers, then train the model to answer the former and abstain from the latter (Zhang et al., 2024; Cheng et al., 2024). To avoid reliance on ground-truth labels, other approaches leverage uncertainty metrics such as semantic entropy to partition the training set into known and unknown questions (Zhang et al., 2024; Tjandra et al., 2024; Xue et al., 2025). However, these methods rely on coarse, global signals, such as a single correctness label or an aggregated uncertainty score, to guide abstention decisions. These approaches don’t account for the nuanced confidence levels of individual samples, resulting in an ambiguous decision boundary. Consequently, models may either excessively abstain from answerable questions or fail to abstain when facing low-confidence inputs, which prevents the model from balancing truthfulness and helpfulness.

To overcome these limitations, we propose **Fine-grained Semantic Confidence Reward (FISCORE, /ˈfɪs.kɔːr/)**, a novel reinforcement learning framework that replaces the conventional global uncertainty reward with a fine-grained, **per-sample** confidence reward. The core insight is that an answer’s intrinsic confidence is reflected in its semantic consensus among all samples, and such confidence can be fully utilized for factual alignment. Although existing methods also sample

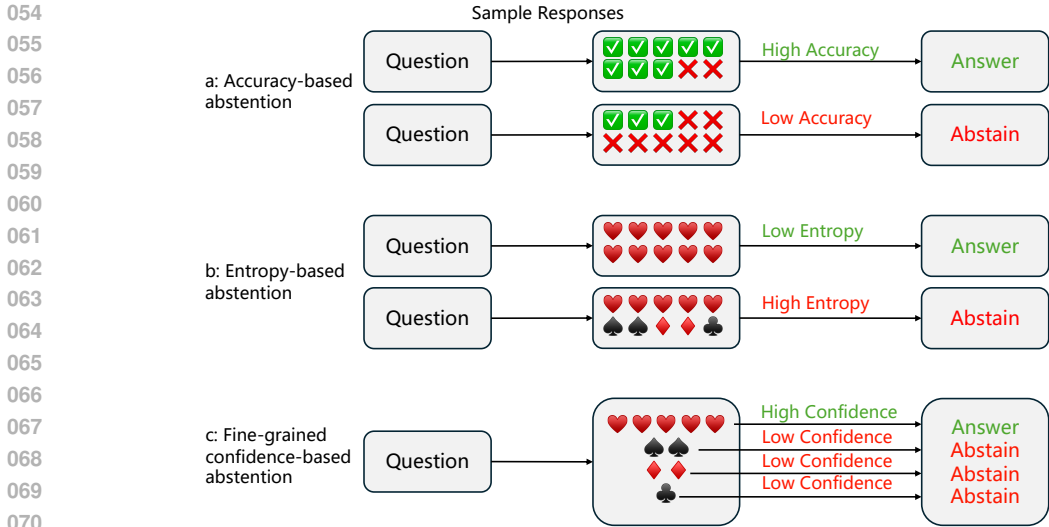


Figure 1: (a) Accuracy-based abstention. (b) Entropy-based abstention. (c) Ours: fine-grained confidence-based abstention.

multiple answers, they typically distill them into a single aggregated uncertainty score. In contrast, as shown in Figure 1, our approach rewards the model for accurately assessing the confidence of each sample. Specifically, we first generate multiple candidate responses and group them into clusters based on their semantic equivalence. We use the cluster size as a proxy for the intrinsic confidence of its member samples: a large cluster implies high confidence, whereas a small cluster implies low confidence. A positive reward is granted only when the model’s expressed confidence (e.g., “sure” or “unsure”) aligns correctly with this intrinsic confidence level. This fine-grained signal incentivizes the model to align its verbalized confidence with semantic consensus, thus enabling it to learn a more nuanced knowledge boundary and make more accurate abstention decisions.

Existing studies indicate that helpfulness and truthfulness are like two ends of a seesaw — enhancing one often leads to a decline in the other (Zhu et al., 2025). We notice that current metrics fail to accurately reflect both aspects, making them inadequate for assessing abstention methods. Regarding this, we propose a reliability metric grounded in the model’s self-awareness of its knowledge boundaries. This metric synthesizes helpfulness and truthfulness, aiming to prevent models from over-rejecting queries or generating hallucinations.

Our main contributions are as follows:

- We design a fine-grained semantic confidence reward (see Section 3.3), which significantly improves the reliability of LLM.
- We propose a new metric (see Section 2.2) that integrates the helpfulness and truthfulness of LLM into a single reliability metric, offering a comprehensive assessment of its awareness of its own knowledge boundaries.
- We experiment with various LLMs and test them on both in-distribution and out-of-distribution data. The results demonstrate that our approach outperforms baselines in most cases, particularly on out-of-distribution data, verifying the effectiveness and generalization capability of the proposed fine-grained semantic confidence reward.

2 PROBLEM FORMULATION

2.1 ABSTENTION FINE-TUNING

To mitigate the risk of hallucination inherent in LLMs, a primary alignment strategy is to fine-tune models to selectively abstain from answering questions that fall beyond their knowledge boundaries (Wen et al., 2025; Li et al., 2025a). The objective is to produce a model that is both helpful, by

Table 1: Abstention confusion matrix.

Initial LLM \ Refined LLM	Correctly answered	Incorrectly answered	Abstained
Known Questions	N_1	N_2	N_3
Unknown Questions	–	N_4	N_5

correctly answering questions within its knowledge boundary, and truthful, by abstaining from those outside it. To formally measure progress towards this goal, we establish an evaluation framework that assesses the performance of a refined LLM relative to the initial LLM before fine-tuning. As illustrated in the abstention confusion matrix in Table 1, we partition the evaluation benchmark into two distinct categories based on the performance of the initial LLM:

- **Known Questions:** The set of questions that the initial LLM answers correctly.
- **Unknown Questions:** The set of questions that the initial LLM answers incorrectly.

The refined LLM is then evaluated on its ability to preserve accuracy on known questions while learning to abstain from unknown questions.

2.2 RELIABILITY EVALUATION

2.2.1 HELPFULNESS AND TRUTHFULNESS

Helpfulness evaluates a model’s capacity to preserve accuracy on questions within its known scope, preventing unnecessary abstention and forgetting. Truthfulness measures its ability to abstain from questions outside that scope, preventing fabricated responses. Following Kim et al. (2024); Feng et al. (2024), we adopt two fine-grained F1 metrics $F1_{ans}$ for helpfulness evaluation and $F1_{abs}$ for truthfulness evaluation:

$$F1_{ans} = \frac{2 \cdot \text{Precision}_{ans} \cdot \text{Recall}_{ans}}{\text{Precision}_{ans} + \text{Recall}_{ans}} = \frac{2N_1}{2N_1 + 2N_2 + N_3 + N_4}, \quad (1)$$

$$\text{Precision}_{ans} = \frac{N_1}{N_1 + N_2 + N_4}, \quad \text{Recall}_{ans} = \frac{N_1}{N_1 + N_2 + N_3}, \quad (2)$$

$$F1_{abs} = \frac{2 \cdot \text{Precision}_{abs} \cdot \text{Recall}_{abs}}{\text{Precision}_{abs} + \text{Recall}_{abs}} = \frac{2N_5}{N_3 + N_4 + 2N_5}, \quad (3)$$

$$\text{Precision}_{abs} = \frac{N_5}{N_3 + N_5}, \quad \text{Recall}_{abs} = \frac{N_5}{N_4 + N_5}. \quad (4)$$

2.2.2 RELIABILITY METRIC

Since increasing truthfulness often leads to a decrease in helpfulness, using either metric alone introduces bias. Therefore, we propose $F1_{rel}$, defined as the harmonic mean of $F1_{ans}$ and $F1_{abs}$, which provides a holistic measure of reliability:

$$F1_{rel} = \frac{2 \cdot F1_{ans} \cdot F1_{abs}}{F1_{ans} + F1_{abs}} = \frac{4N_1N_5}{4N_1N_5 + 2N_2N_5 + N_1N_3 + N_1N_4 + N_3N_5 + N_4N_5}. \quad (5)$$

The $F1_{rel}$ metric holistically considers all categories in the abstention confusion matrix in Table 1 and exhibits the following desiderata: it monotonically increases with the correctly answered known questions (N_1) and correct abstention from unknown questions (N_5); it decreases monotonically with both types of error: incorrect answer (N_2, N_3) and over-abstention (N_4), and therefore is intuitively reasonable. In the ideal case ($N_2 = N_3 = N_4 = 0, N_1 + N_5 = N$), $F1_{rel} = 1$; in the worst case ($N_1 = N_5 = 0, N_2 + N_3 + N_4 = N$), $F1_{rel} = 0$, where N is the number of all questions. This metric effectively balances helpfulness and truthfulness, mitigating the risks of overly conservative or aggressive model behavior. A recent metric, reliability score (RS) (Xu et al., 2024a), has been proposed to achieve similar balancing effect, but we show that it is flawed in encouraging hallucination (see Appendix C.1).

3 METHOD

Inspired by the powerful generalization capabilities of reinforcement learning (Chu et al., 2025) and its success in reasoning (Guo et al., 2025; Zuo et al., 2025), we incorporate abstention decisions into the rollout process and train LLMs to abstain appropriately via fine-grained confidence reward (FISCORE). In this section, we first show the details of training FISCORE and formalize the reinforcement learning objective. Then, we demonstrate the prompt template that guides the LLMs to generate rollout in the defined format. Finally, we introduce the reward modeling to guide the optimization of reinforcement learning to improve reliability.

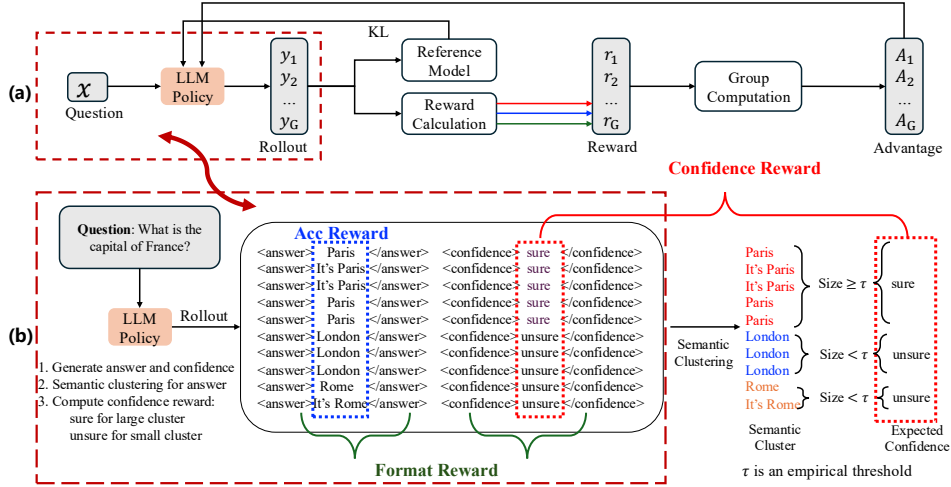


Figure 2: The training overview of our method. (a) The GRPO pipeline. (b) A detailed example of how FISCORE works.

3.1 GROUP RELATIVE POLICY OPTIMIZATION

We reward the model based on the sample-specific semantic confidence among a group of its own sampled responses, which requires a reinforcement learning algorithm capable of handling multiple samples per prompt. Group Relative Policy Optimization (GRPO) (Shao et al., 2024) naturally fits this setting: it samples a group of G outputs (or *rollouts*) for each question x , which directly supports our semantic clustering-based confidence reward computation. Formally, for each question x , GRPO samples a group of G outputs $Y = \{y_1, y_2, \dots, y_G\}$ from the old policy $\pi_{\theta_{\text{old}}}$ and optimizes the new policy π_{θ} by maximizing the following objective:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{[x \sim P(Q), \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(Y|x)]} \frac{1}{G} \sum_{i=1}^G \left(\min \left(\frac{\pi_{\theta}(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)} A_i, \text{clip} \left(\frac{\pi_{\theta}(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)}, 1 - \epsilon, 1 + \epsilon \right) A_i \right) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} || \pi_{\text{ref}}) \right), \quad (6)$$

where $A_i = (r_i - \text{mean}(\{r_i\}_{i=1}^G)) / \text{std}(\{r_i\}_{i=1}^G)$ is the normalized advantage of the i th rollout in the current group, ϵ is the clipping ratio, and β is the KL loss coefficient. Moreover, the objective also incorporates a KL divergence penalty, specifically to prevent excessive divergence of the policy from the original reference policy model. The GRPO pipeline is shown in Figure 2 (a).

3.2 TRAINING TEMPLATE

As we describe the confidence reward in Figure 2 (b), it is essential for the policy LLM to generate responses that conform to the specified format. To facilitate this, we design a prompt template to guide the policy model in producing an answer along with an associated confidence level. As shown in Table 2, the model is instructed to format its responses into two distinct parts. First, the model

Table 2: Training template. The question is appended at the end during training and inference.

You are a helpful and truthful AI Assistant. Given a question, you should provide responses that include answer and confidence. You first answer the question as briefly as possible enclosed by `<answer>` and `</answer>`, and then provide your confidence in sure or unsure about the answer, enclosed by `<confidence>` and `</confidence>`. Respond in the following format:

```
<answer> ... </answer>
<confidence> sure or unsure </confidence>
```

Question:

provides the answer wrapped in the `<answer>` `</answer>` tag. Second, the model expresses a binary confidence level (sure or unsure) wrapped in the `<confidence>` `</confidence>` tag. This `<answer>` + `<confidence>` structure enforces the model to reflect on its confidence after providing an answer. In this work, we interpret an unsure confidence as an act of abstention.

3.3 REWARDS FUNCTION DESIGN

Figure 2 (b) illustrates the proposed FiSCORE. We first partition the responses using a proxy model that captures semantic equivalence, then introduce the confidence reward and auxiliary rewards.

3.3.1 FINE-GRAINED SEMANTIC CONFIDENCE REWARD

Some studies employ reinforcement learning with confidence rewards to enhance the reasoning performance of LLMs (Zuo et al., 2025; Shafayat et al., 2025). We extend this idea to the abstention domain, which guides LLMs to answer when confident and abstain otherwise.

Semantic Clustering. To compute the semantic confidence reward for individual rollouts during GRPO training, we initially organize responses that express equivalent meanings through semantic clustering (Kuhn et al., 2023). These semantic clusters represent equivalence classes derived from a semantic equivalence relation that satisfies reflexivity, symmetry, and transitivity properties while capturing the semantic similarity between textual inputs. The practical implementation of equivalence via bidirectional entailment predictions generated by a Natural Language Inference (NLI) model, specifically DeBERTa (He et al., 2021), which determines the relationships between text pairs as either entailment, neutral, or contradiction. Texts are considered semantically equivalent when they mutually entail one another in both directions. We follow Kuhn et al. (2023) employing a greedy algorithm that assigns each response to an existing semantic cluster if it demonstrates semantic equivalence with any member of that cluster.

Semantic Confidence. The formulation of FiSCORE is based on the intuition that generative semantic consensus reflects model confidence. We assume that a larger semantic cluster implies strong consensus among the generated responses, reflecting higher confidence. Conversely, smaller or singleton clusters suggest divergent or less-supported meanings, thus implying lower confidence. Following this principle, we quantify the confidence of each rollout using a measure derived from the cardinality of the semantic cluster it belongs to. This per-sample approach distinguishes our method from previous work that relies on a single, global metric like semantic entropy calculated over all responses. For a given rollout y_i and its corresponding semantic cluster C_i , the confidence reward function is defined as:

$$R_c = \begin{cases} 1, & \text{if } (|C_i| \geq \tau \text{ and } U = \text{“sure”}) \text{ or } (|C_i| < \tau \text{ and } U = \text{“unsure”}) \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

Here, τ is an empirical threshold of abstention, which we set to $\lceil \frac{G}{2} \rceil$ in all experiments, where G is the total number of rollouts. U is the model’s expressed confidence, either “sure” or “unsure”, extracted from the `<confidence>` `</confidence>` tag within the rollout. By incentivizing this alignment, the reward function effectively teaches the model a reliable abstention policy: to answer when semantic confidence is high and to abstain otherwise.

3.3.2 AUXILIARY REWARDS

Accuracy Reward. To prevent the model from engaging in reward hacking by learning to maximize the confidence reward on a uniform but incorrect answer, we introduce an auxiliary accuracy reward. This reward is granted only when the answer is semantically equivalent to the ground truth solution.

$$R_a = \begin{cases} 1, & \text{if the answer and the solution are semantically equivalent} \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

where the answer is extracted between `<answer>` `</answer>` tags.

Format Reward. We encourage the model to enclose the answer within `<answer>` `</answer>` tags and to present the confidence enclosed within `<confidence>` `</confidence>` tags. The format reward R_f is composed of two components. For each correctly used tag (`<answer>`, `</answer>`, `<confidence>`, `</confidence>`) that appears in the rollout, we assign a reward of 0.125. To allow the model to determine its confidence based on its answer, it is necessary to place the answer before the confidence expression. We assign an additional reward of 0.5 if the tags appear in the correct order.

The total reward R_{total} is defined as a weighted sum of three distinct components:

$$R_{total} = w_c R_c + w_a R_a + w_f R_f, \quad (9)$$

where w_c , w_a , w_f are the respective weights for each reward component. Notably, only when the format is correct ($R_f = 1$) will the other two rewards be added.

4 EXPERIMENTS

4.1 EXPERIMENT SETTING

Datasets. We use four datasets of question-answering (QA) task for evaluation, including **Pararel** (Elazar et al., 2021) which contains cloze-style questions of relation prediction, **TriviaQA** (Joshi et al., 2017) which contains general knowledge QA pairs, **Natural Questions (NQ)** (Kwiatkowski et al., 2019) which contains questions from aggregated queries to Google Search, and **SciQ** (Welbl et al., 2017) which contains science exam questions. We use the **Pararel** dataset for training.

Baselines. We compare FISCORE with three types of baselines.

1. **Prompting method:** In-Context Learning (ICL), In-Context Learning with refusal examples (ICL-IDK), and In-Context Learning with confidence expression (ICL-Unsure). We take the performance of ICL as the performance of the initial model.
2. **Supervised fine-tuning methods:** R-Tuning (Zhang et al., 2024), R-Tuning-U (Zhang et al., 2024), and SE-Tuning (Tjandra et al., 2024).
3. **RL-based method:** GRPO with semantic entropy (GRPO-SE), which adopts the training data as SE-Tuning and is built on GRPO.

Evaluation. We use four metrics: $F1_{ans}$, $F1_{abs}$, $F1_{rel}$, and accuracy (Acc) for evaluation, which are detailed in Section 2.2. We evaluate the correctness of the model’s response by prompting with 6-shot examples. Following Xue et al. (2025), we adopt bidirectional string matching to evaluate answer correctness. Here, we treat a response as an abstention if the confidence part is unsure. More details about the evaluation are listed in Appendix D.

Implementation Details. We choose Llama3-8B-Instruct (Grattafiori et al., 2024) and Qwen2.5-7B-Instruct (Yang et al., 2024) in our experiments. The reinforcement learning framework is built on Open-R1 (Face, 2025). We set the reward weights $w_c = 1$, $w_a = 4$ and $w_f = 2$. We use Deberta-v2-xlarge-mnli (He et al., 2021) as the NLI model for semantic clustering. The sampling number is set to 10, and the threshold τ is set to 5. All experiments are implemented on Nvidia L40-48GB GPUs. The training epoch is set to 1. We set the sampling temperature to 1.0 during the GRPO training. During inference, we utilize the vLLM framework to accelerate the process and employ a greedy search strategy to generate responses. Hyperparameters and additional configurations are detailed in Appendix B.

Table 3: Performance on in-domain and out-of-domain knowledge-intensive QA datasets. All results are multiplied by 100. Notably, ICL is the initial model, so we do not calculate its F1 scores.

Method	Pararel (ID)				TriviaQA (OOD)				NQ (OOD)				SciQ (OOD)			
	F1 _{ans}	F1 _{abs}	F1 _{rel}	Acc	F1 _{ans}	F1 _{abs}	F1 _{rel}	Acc	F1 _{ans}	F1 _{abs}	F1 _{rel}	Acc	F1 _{ans}	F1 _{abs}	F1 _{rel}	Acc
Llama3-8B-Instruct																
ICL	-	-	-	48.8	-	-	-	68.4	-	-	-	37.1	-	-	-	65.3
ICL-IDK	62.1	64.3	63.2	31.7	69.8	32.6	44.5	55.3	43.5	67.9	53.0	18.1	71.8	43.9	54.5	56.4
ICL-Unsure	70.4	68.2	69.3	41.9	67.7	46.4	55.0	51.7	45.9	18.0	25.9	30.4	64.7	17.9	28.0	53.1
R-Tuning	78.3	79.6	79.0	40.5	47.7	51.5	49.5	26.3	31.2	72.0	43.5	10.8	32.1	51.4	39.5	14.6
R-Tuning-U	73.8	76.6	75.2	40.5	42.6	53.2	47.3	20.4	31.8	74.0	44.5	10.3	49.9	49.7	49.8	28.0
SE-Tuning	78.2	75.0	76.6	45.9	44.9	50.7	47.7	24.0	27.7	72.2	40.0	9.4	22.4	50.4	31.0	9.2
GRPO-SE	70.9	77.3	73.9	35.1	70.5	60.2	65.0	47.3	37.2	73.9	49.5	13.0	64.4	51.2	57.0	46.4
FISCoRE	83.3	67.1	74.3	59.4	82.6	62.0	70.8	61.7	53.1	72.5	61.3	21.7	74.9	53.3	62.3	52.4
Qwen2.5-7B-Instruct																
ICL	-	-	-	44.1	-	-	-	51.7	-	-	-	27.6	-	-	-	68.8
ICL-IDK	60.1	68.6	64.1	27.4	66.6	49.3	56.6	44.9	38.1	50.1	43.3	19.4	71.6	14.8	24.5	60.5
ICL-Unsure	77.7	82.9	80.2	35.8	71.2	51.7	59.9	49.8	40.2	39.7	39.9	21.9	71.4	37.6	49.2	57.3
R-Tuning	79.3	86.4	82.7	31.8	44.0	71.0	54.4	15.1	25.1	84.5	38.7	4.5	43.8	53.9	48.3	19.9
R-Tuning-U	62.1	80.5	70.1	21.3	40.1	69.8	51.0	13.6	32.3	84.8	46.7	6.3	43.9	52.5	47.8	20.4
SE-Tuning	71.4	82.3	76.5	28.0	33.8	69.3	45.4	10.7	32.8	84.2	47.2	6.7	55.4	52.7	54.0	34.5
GRPO-SE	75.2	83.3	79.0	31.7	71.6	72.4	72.0	41.6	42.8	78.0	55.3	14.2	76.7	50.7	61.0	57.3
FISCoRE	80.8	72.9	76.6	51.9	83.3	79.1	81.1	51.3	57.3	80.1	66.8	20.4	79.2	54.5	64.6	58.0

4.2 MAIN RESULTS

We present the main experimental results of FISCoRE and all baseline methods in Table 3, and we show the methods based on different LLMs. The results reveal several key findings:

Generalizable Reliability. We focus on metric $F1_{rel}$ that assesses a model’s overall reliability by balancing helpfulness ($F1_{ans}$) and truthfulness ($F1_{abs}$). A key finding is that the high reliability demonstrated by some baselines on in-domain (ID) data is brittle and does not generalize. For example, R-Tuning using the Llama model achieves a high $F1_{rel}$ score of 79.0 on the Pararel (ID) dataset by directly learning to abstain from unknown questions. However, this specialized performance collapses on out-of-distribution (OOD) tasks, with its score dropping to just 49.5 of $F1_{rel}$ in TriviaQA. In contrast, the proposed FISCoRE demonstrates a more robust, reliable performance that excels at generalization. It achieves this by attaining a superior balance between answering known questions ($F1_{ans}$ of 83.3) and correctly abstaining ($F1_{abs}$ of 67.1). This balanced approach leads to a high $F1_{rel}$ score of 74.3 on the ID dataset, which shows a much smaller drop to 70.8 on the OOD TriviaQA dataset. This consistent outperformance on all three OOD datasets proves that the knowledge boundary learned by our approach is more generalizable. Compared to SE-Tuning, GRPO-SE shows stronger generalization: SE-Tuning performs well on ID data but degrades significantly on OOD sets, while GRPO-SE maintains robust performance across both, highlighting the advantage of reinforcement learning over supervised fine-tuning.

Effectiveness of the Fine-grained Semantic Confidence Reward. We compare our approach directly with GRPO-SE, a baseline employing the same GRPO algorithm but utilizing a coarse semantic uncertainty reward. The results demonstrate a clear advantage: our method FISCoRE significantly outperforms GRPO-SE across all OOD datasets and different LLMs in terms of $F1_{rel}$. Given that the primary distinction between the two methods is the reward function, this comparison provides empirical evidence that our fine-grained confidence reward is substantially more effective than a coarse, global entropy-based signal.

Comparison between Different LLMs. We fine-tuned both the Llama3-8B-Instruct and Qwen2.5-7B-Instruct models on the Pararel dataset using our FISCoRE method alongside established baseline abstention techniques. From the results on TriviaQA and NQ we can see that although Llama

achieved higher raw accuracy than Qwen both before and after fine-tuning, its overall reliability (measured by $F1_{rel}$) remained lower. This disparity indicates that a model’s extensive knowledge capacity does not ensure precise self-awareness of its knowledge boundaries. It also highlights that accuracy alone provides an inadequate proxy for gauging improvements in abstention alignment.

4.3 ABLATION STUDY




Figure 3: (a) Experiments of $F1_{rel}$ and accuracy of various sampling number G on TriviaQA on Qwen2.5-7B-Instruct. (b) Experiments of $F1_{ans}$, $F1_{abs}$ and $F1_{rel}$ of various accuracy reward weight w_a on TriviaQA on Qwen2.5-7B-Instruct.

Impact of Sampling Number. We investigate the impact of the sampling number G on performance in Figure 3 (a). The results demonstrate that for our method FiSCORE, both reliability ($F1_{rel}$) and accuracy improve as G increases, consistently outperforming the GRPO-SE baseline. Both metrics reach their optimal or near-optimal levels at $G = 10$. Since performance gains become negligible beyond this point, we select $G = 10$ as the default setting for all experiments to maximize performance while maintaining computational efficiency.

Impact of Accuracy Reward Weight. We evaluate the impact of the accuracy reward weight w_a while fixing the confidence reward weight w_c at 1. When w_a is set to 0, the model achieves a high $F1_{abs}$ but a near-zero $F1_{ans}$, indicating that the model abstains from almost all questions. This occurs because the model receives rewards for generating low-confidence incorrect answers, a phenomenon known as reward hacking. Shafayat et al. (2025) demonstrate similar results: relying on self-confidence as the only reward leads to model collapse in the later stages of training. Therefore, it is necessary to combine accuracy reward and confidence reward for training. Introducing the accuracy reward improves $F1_{rel}$. As the accuracy reward weight increases, both $F1_{ans}$ and $F1_{rel}$ rise steadily. Notably, weights greater than 1 produce higher $F1_{rel}$ scores, likely because obtaining accuracy rewards is more difficult than obtaining confidence rewards, thus requiring greater emphasis.

Table 4: Ablation on threshold.

	$F1_{ans}$	$F1_{abs}$	$F1_{rel}$	Acc
3	79.3	65.6	71.8	54.7
4	82.4	75.4	78.8	53.0
5	83.3	79.1	81.1	51.3
6	82.1	78.3	80.1	49.9
7	74.7	81.2	77.8	36.5

Table 5: Ablation on reward functions.

	$F1_{ans}$	$F1_{abs}$	$F1_{rel}$	Acc
FiSCORE	83.3	79.1	81.1	51.3
w/o Format reward	71.5	43.6	54.2	51.5
w/o Confidence reward	69.5	2.1	4.1	56.1
w/ Continuous reward	83.0	82.9	82.9	49.1

Impact of Abstention Threshold. We conduct ablation experiments on the TriviaQA dataset using the Qwen2.5-7B-Instruct model to investigate the impact of the abstention threshold τ . As shown in the Table 4, $F1_{abs}$ increases with rising τ , while $F1_{ans}$ and Acc exhibit a downward trend. At $\tau = 5$, truthfulness and helpfulness achieve a balanced level, yielding the maximum $F1_{rel}$. Therefore, different thresholds can be employed for training across scenarios with varying tolerance levels for hallucinations.

Ablation on Reward Functions. Table 5 presents the results of reward function ablation experiments conducted on the TriviaQA dataset using the Qwen2.5-7B-Instruct model. We first validate the necessity of the format reward. We remove the format reward and observe a decrease in $F1_{ans}$, a significant drop in $F1_{abs}$, while accuracy remains almost unaffected. The reason lies in semantic clustering: we extract answers within the answer tags using rule-based methods. The lack of format reward leads to inaccurate tags, which degrade semantic clustering performance and ultimately increase confidence bias. To validate the effectiveness of the proposed fine-grained semantic confidence reward, we removed this reward and observed that while accuracy remained stable, $F1_{abs}$ dropped to near zero. This demonstrates that the model possesses a certain degree of confidence perception capability, and our confidence reward can significantly enhance this ability. Finally, we train Qwen2.5-7B-Instruct with continuous confidence rewards, where the model samples 10 candidate answers and predicts confidence scores scaled from 10% to 100%. We find that continuous rewards achieve better reliability than binary discrete rewards. Continuous rewards do not require retraining when adapting the abstention threshold, resulting in lower applied costs.

4.4 DETAILED ANALYSIS




Figure 4: Percentage of prediction types among different methods. We choose SE-Tuning, GRPO-SE, and FiSCORE.

Figure 4 compares the performance of three methods: SE-Tuning, GRPO-SE, and FiSCORE. The primary goal is to maximize correctly answered known questions (N_1) and abstained unknown questions (N_5), while minimizing errors (N_2, N_4) and unnecessary abstentions (N_3). Note that Pararel is the ID dataset, whereas all the others are OOD. On OOD datasets, SE-Tuning tends to over-abstain, reflected by a larger N_3 , indicating a conservative strategy that suppresses hallucinations at the cost of unnecessary refusals on answerable cases. By contrast, GRPO-SE and FiSCORE show higher hallucination rates than SE-Tuning, which reflects the seesaw effect between helpfulness and truthfulness: reducing over-abstention encourages the model to answer more often and thus increases the chance of hallucinations relative to an overly conservative baseline. Importantly, in most cases FiSCORE attains a higher N_1+N_5 than the baselines, indicating a better overall balance between answering known questions and abstaining on unknown ones, and hence higher reliability.

5 RELATED WORK

Uncertainty Estimation and Confidence Elicitation. Uncertainty estimation aims to quantify this risk before LLMs respond to a request, which is an effective tool for detecting and mitigating hallucinations of LLMs (Xia et al., 2025). Likelihood-based methods estimate uncertainty by computing probabilities for salient tokens, capturing the model’s confidence in specific outputs (Duan et al., 2024; Lin et al., 2024). Sampling-based methods generate multiple candidate responses, cluster them, and quantify uncertainty based on the degree of inter-response consistency (Kuhn et al., 2023; Qiu & Miikkulainen, 2024; Nikitin et al., 2024; Aichberger et al., 2025). Probing-based methods leverage internal model representations to train classifiers that predict uncertainty or detect hallucinations, often without requiring additional generation (Han et al., 2024; Marks & Tegmark, 2024;

Gottesman & Geva, 2024; Azaria & Mitchell, 2023). Confidence elicitation aims to elicit the confidence of LLMs in their answers (Geng et al., 2024). Prompting methods directly prompt LLMs to express confidence (Lin et al., 2022; Kadavath et al., 2022). Recent studies (Xu et al., 2024b; Damani et al., 2025) employ chain-of-thought (CoT) training to teach LLMs to reason about their confidence. Zheng et al. (2025) adopt binary confidence elicitation (“sure” vs. “unsure”) as a proxy for abstention, enabling models to discard low-confidence answers post hoc. In this work, we introduce a fine-grained semantic confidence reward to cultivate a more nuanced self-awareness of the model’s confidence.

Abstention Fine-Tuning. Many studies explore fine-tuning LLMs to abstain from answering questions beyond their knowledge boundaries (Wen et al., 2025). Brahman et al. (2024); Cheng et al. (2024); Xu et al. (2024a) propose learning from preferences via direct preference optimization (Rafailov et al., 2023) to train LLMs to abstain from answering unknown questions. Li et al. (2025b) introduce adaptive contrastive preference learning to calibrate abstention behavior. Zhang et al. (2024); Tjandra et al. (2024); Xue et al. (2025) leverage uncertainty to construct abstention-aware datasets, replacing the ground-truth labels of uncertain questions with “I don’t know”. Yang et al. (2025); Deng et al. (2024) train LLMs to reason about their own uncertainty and provide post-hoc explanations for why a given question is unanswerable. Huang et al. (2025); Cohen et al. (2024) introduce a special “rejection” token into the model’s vocabulary and design an objective function that reallocates probability mass toward the “rejection” token when the model encounters uncertain predictions. In this work, we go beyond aggregated uncertainty by utilizing fine-grained, per-sample confidence to sharpen the model’s awareness of its knowledge limits.

6 CONCLUSION

In this paper, we focus on enhancing the reliability of LLMs, which we define as the ability to accurately answer questions within their knowledge boundaries and abstain from those beyond them. First, we propose a new comprehensive reliability metric by calculating the harmonic mean of the F1 scores for answering and abstaining. This metric integrates model helpfulness and truthfulness into a single, robust score that is monotonically sensitive to all error categories. We then introduce FiSCORE, a reinforcement learning method with a fine-grained semantic confidence reward. This reward mechanism, derived from the semantic confidence among multiple generated responses, guides the model to retain the answer only when generative consensus is high and to discard it when it is low, thereby enabling the model to make abstention decisions more aligned with its knowledge boundaries. Our experiments on both in-domain and out-of-distribution datasets show that FiSCORE achieves a better balance between helpfulness and truthfulness than baselines, thereby improving overall model reliability.

REPRODUCIBILITY STATEMENT

To ensure reproducibility of our work, we provide detailed instructions and the necessary code. The source code for training and evaluation is available in the supplementary material. We also provide the baseline code.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Lukas Aichberger, Kajetan Schweighofer, Mykyta Ielanskyi, and Sepp Hochreiter. Improving uncertainty estimation through semantically diverse language generation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=HSi4VetQLj>.
- Amos Azaria and Tom Mitchell. The internal state of an LLM knows when it’s lying. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. URL <https://openreview.net/forum?id=y2V6YgLaW7>.

- 540 Faeze Brahman, Sachin Kumar, Vidhisha Balachandran, Pradeep Dasigi, Valentina Pyatkin, Abhi-
541 lasha Ravichander, Sarah Wiegrefe, Nouha Dziri, Khyathi Chandu, Jack Hessel, Yulia Tsvetkov,
542 Noah A. Smith, Yejin Choi, and Hannaneh Hajishirzi. The art of saying no: Contextual noncom-
543 pliance in language models. In *The Thirty-eight Conference on Neural Information Processing*
544 *Systems Datasets and Benchmarks Track*, 2024. URL [https://openreview.net/forum?](https://openreview.net/forum?id=f1UL4wNlw6)
545 [id=f1UL4wNlw6](https://openreview.net/forum?id=f1UL4wNlw6).
- 546 Qinyuan Cheng, Tianxiang Sun, Xiangyang Liu, Wenwei Zhang, Zhangyue Yin, Shimin Li,
547 Linyang Li, Zhengfu He, Kai Chen, and Xipeng Qiu. Can AI assistants know what they
548 don’t know? In *Forty-first International Conference on Machine Learning*, 2024. URL
549 <https://openreview.net/forum?id=girxGkdECL>.
- 550
551 Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V.
552 Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation
553 model post-training, 2025. URL <https://arxiv.org/abs/2501.17161>.
- 554 Roi Cohen, Konstantin Dobler, Eden Biran, and Gerard de Melo. I don’t know: Ex-
555 plicit modeling of uncertainty with an [idk] token. In A. Globerson, L. Mackey, D. Bel-
556 grave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural In-*
557 *formation Processing Systems*, volume 37, pp. 10935–10958. Curran Associates, Inc.,
558 2024. URL [https://proceedings.neurips.cc/paper_files/paper/2024/](https://proceedings.neurips.cc/paper_files/paper/2024/file/14c018d2e72c521605b0567029ef0efb-Paper-Conference.pdf)
559 [file/14c018d2e72c521605b0567029ef0efb-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/14c018d2e72c521605b0567029ef0efb-Paper-Conference.pdf).
- 560 Mehul Damani, Isha Puri, Stewart Slocum, Idan Shenfeld, Leshem Choshen, Yoon Kim, and Jacob
561 Andreas. Beyond binary rewards: Training lms to reason about their uncertainty. *arXiv preprint*
562 *arXiv:2507.16806*, 2025.
- 563
564 Yang Deng, Yong Zhao, Moxin Li, See-Kiong Ng, and Tat-Seng Chua. Don’t just say “I don’t
565 know”! self-aligning large language models for responding to unknown questions with expla-
566 nations. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the*
567 *2024 Conference on Empirical Methods in Natural Language Processing*, pp. 13652–13673, Mi-
568 ami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/
569 [v1/2024.emnlp-main.757](https://aclanthology.org/2024.emnlp-main.757). URL [https://aclanthology.org/2024.emnlp-main.](https://aclanthology.org/2024.emnlp-main.757/)
570 [757/](https://aclanthology.org/2024.emnlp-main.757/).
- 571 Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura,
572 and Kaidi Xu. Shifting attention to relevance: Towards the predictive uncertainty quantifica-
573 tion of free-form large language models. In Lun-Wei Ku, Andre Martins, and Vivek Sriku-
574 mar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Lin-*
575 *guistics (Volume 1: Long Papers)*, pp. 5050–5063, Bangkok, Thailand, August 2024. Asso-
576 ciation for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.276. URL [https:](https://aclanthology.org/2024.acl-long.276/)
577 [//aclanthology.org/2024.acl-long.276/](https://aclanthology.org/2024.acl-long.276/).
- 578 Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich
579 Schütze, and Yoav Goldberg. Measuring and improving consistency in pretrained language mod-
580 els. *Transactions of the Association for Computational Linguistics*, 9:1012–1031, 2021.
- 581 Hugging Face. Open rl: A fully open reproduction of deepseek-rl, January 2025. URL [https:](https://github.com/huggingface/open-rl)
582 [//github.com/huggingface/open-rl](https://github.com/huggingface/open-rl).
- 583
584 Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov.
585 Don’t hallucinate, abstain: Identifying LLM knowledge gaps via multi-LLM collaboration. In
586 Lun-Wei Ku, Andre Martins, and Vivek Sriku-
587 mar (eds.), *Proceedings of the 62nd Annual Meeting*
588 *of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14664–14690,
589 Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/
2024.acl-long.786. URL <https://aclanthology.org/2024.acl-long.786/>.
- 590 Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. A survey
591 of confidence estimation and calibration in large language models. In Kevin Duh, Helena Gomez,
592 and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter*
593 *of the Association for Computational Linguistics: Human Language Technologies (Volume 1:*
Long Papers), pp. 6577–6595, Mexico City, Mexico, June 2024. Association for Computational

- 594 Linguistics. doi: 10.18653/v1/2024.naacl-long.366. URL [https://aclanthology.org/](https://aclanthology.org/2024.naacl-long.366/)
595 2024.naacl-long.366/.
596
- 597 Daniela Gottesman and Mor Geva. Estimating knowledge in large language models without gener-
598 ating a single token. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-*
599 *ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.
600 3994–4019, Miami, Florida, USA, November 2024. Association for Computational Linguis-
601 tics. doi: 10.18653/v1/2024.emnlp-main.232. URL [https://aclanthology.org/2024.](https://aclanthology.org/2024.emnlp-main.232/)
602 [emnlp-main.232/](https://aclanthology.org/2024.emnlp-main.232/).
- 603 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
604 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd
605 of models. *arXiv preprint arXiv:2407.21783*, 2024.
- 606
607 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
608 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
609 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- 610 Jiatong Han, Jannik Kossen, Muhammed Razzak, Lisa Schut, Shreshth A Malik, and Yarin Gal. Se-
611 mantic entropy probes: Robust and cheap hallucination detection in LLMs. In *ICML 2024 Work-*
612 *shop on Foundation Models in the Wild*, 2024. URL [https://openreview.net/forum?](https://openreview.net/forum?id=Zd0XLR6JKn)
613 [id=Zd0XLR6JKn](https://openreview.net/forum?id=Zd0XLR6JKn).
- 614 Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced bert
615 with disentangled attention. In *International Conference on Learning Representations*, 2021.
616 URL <https://openreview.net/forum?id=XPZTaotutsD>.
617
- 618 Lei Huang, Xiaocheng Feng, Weitao Ma, Yuchun Fan, Xiachong Feng, Yuxuan Gu, Yangfan Ye,
619 Liang Zhao, Weihong Zhong, Baoxin Wang, Dayong Wu, Guoping Hu, Lingpeng Kong, Tong
620 Xiao, Ting Liu, and Bing Qin. Alleviating hallucinations from knowledge misalignment in large
621 language models via selective abstention learning. In Wanxiang Che, Joyce Nabende, Ekate-
622 rina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting*
623 *of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 24564–24579,
624 Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-
625 251-0. doi: 10.18653/v1/2025.acl-long.1199. URL [https://aclanthology.org/2025.](https://aclanthology.org/2025.acl-long.1199/)
626 [acl-long.1199/](https://aclanthology.org/2025.acl-long.1199/).
- 627 Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly
628 supervised challenge dataset for reading comprehension. In Regina Barzilay and Min-Yen Kan
629 (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*
630 *(Volume 1: Long Papers)*, pp. 1601–1611, Vancouver, Canada, July 2017. Association for Com-
631 putational Linguistics. doi: 10.18653/v1/P17-1147. URL [https://aclanthology.org/](https://aclanthology.org/P17-1147/)
632 [P17-1147/](https://aclanthology.org/P17-1147/).
- 633 Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez,
634 Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language mod-
635 els (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.
636
- 637 Hyuhng Joon Kim, Youna Kim, Cheonbok Park, Junyeob Kim, Choonghyun Park, Kang Min Yoo,
638 Sang-goo Lee, and Taeuk Kim. Aligning language models to explicitly handle ambiguity. In
639 Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference*
640 *on Empirical Methods in Natural Language Processing*, pp. 1989–2007, Miami, Florida, USA,
641 November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.
642 119. URL <https://aclanthology.org/2024.emnlp-main.119/>.
- 643 Hyuhng Joon Kim, Youna Kim, Sang-goo Lee, and Taeuk Kim. When to speak, when to abstain:
644 Contrastive decoding with abstention. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and
645 Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for*
646 *Computational Linguistics (Volume 1: Long Papers)*, pp. 9710–9730, Vienna, Austria, July 2025.
647 Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.
[acl-long.479](https://aclanthology.org/2025.acl-long.479/). URL <https://aclanthology.org/2025.acl-long.479/>.

- 648 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances
649 for uncertainty estimation in natural language generation. In *The Eleventh International Confer-*
650 *ence on Learning Representations*, 2023. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=VD-AYtP0dve)
651 [VD-AYtP0dve](https://openreview.net/forum?id=VD-AYtP0dve).
- 652 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris
653 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion
654 Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav
655 Petrov. Natural questions: A benchmark for question answering research. *Transactions of the*
656 *Association for Computational Linguistics*, 7:452–466, 2019. doi: 10.1162/tacl_a.00276. URL
657 <https://aclanthology.org/Q19-1026/>.
- 658 Moxin Li, Yong Zhao, Wenxuan Zhang, Shuaiyi Li, Wenya Xie, See-Kiong Ng, Tat-Seng Chua, and
659 Yang Deng. Knowledge boundary of large language models: A survey. In *Proceedings of the*
660 *63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
661 pp. 5131–5157, Vienna, Austria, July 2025a. Association for Computational Linguistics. ISBN
662 979-8-89176-251-0. URL <https://aclanthology.org/2025.acl-long.256/>.
- 663 Yinghui Li, Haojing Huang, Jiayi Kuang, Yangning Li, Shu-Yu Guo, Chao Qu, Xiaoyu Tan, Hai-
664 Tao Zheng, Ying Shen, and Philip S. Yu. Refine knowledge of large language models via adaptive
665 contrastive learning. In *The Thirteenth International Conference on Learning Representations*,
666 2025b. URL <https://openreview.net/forum?id=HqjRlT65WX>.
- 667 Stephanie Lin, Jacob Hilton, and Owain Evans. Teaching models to express their uncertainty in
668 words. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=8s8K2UZGTZ>.
- 669 Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Contextualized sequence likelihood: Enhanced
670 confidence scores for natural language generation. In Yaser Al-Onaizan, Mohit Bansal, and
671 Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natu-*
672 *ral Language Processing*, pp. 10351–10368, Miami, Florida, USA, November 2024. Associa-
673 tion for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.578. URL <https://aclanthology.org/2024.emnlp-main.578/>.
- 674 Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language
675 model representations of true/false datasets. In *First Conference on Language Modeling*, 2024.
676 URL <https://openreview.net/forum?id=aaJyHYjjsk>.
- 677 Alexander V Nikitin, Jannik Kossen, Yarin Gal, and Pekka Marttinen. Kernel language entropy:
678 Fine-grained uncertainty quantification for LLMs from semantic similarities. In *The Thirty-*
679 *eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=j2wCrWmgMX>.
- 680 Xin Qiu and Risto Miikkulainen. Semantic density: Uncertainty quantification for large language
681 models through confidence measurement in semantic space. In *The Thirty-eighth Annual Confer-*
682 *ence on Neural Information Processing Systems*, 2024. URL [https://openreview.net/](https://openreview.net/forum?id=LOH6qzI7T6)
683 [forum?id=LOH6qzI7T6](https://openreview.net/forum?id=LOH6qzI7T6).
- 684 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
685 Finn. Direct preference optimization: Your language model is secretly a reward model. In
686 A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in*
687 *Neural Information Processing Systems*, volume 36, pp. 53728–53741. Curran Associates, Inc.,
688 2023. URL [https://proceedings.neurips.cc/paper_files/paper/2023/](https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf)
689 [file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf).
- 690 Sheikh Shafayat, Fahim Tajwar, Ruslan Salakhutdinov, Jeff Schneider, and Andrea Zanette. Can
691 large reasoning models self-train?, 2025. URL <https://arxiv.org/abs/2505.21444>.
- 692 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
693 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
694 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- 695
696
697
698
699
700
701

- 702 Benedict Aaron Tjandra, Muhammed Razzak, Jannik Kossen, Kunal Handa, and Yarin Gal. Fine-
703 tuning large language models to appropriately abstain with semantic entropy. In *Neurips Safe*
704 *Generative AI Workshop 2024*, 2024. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=oYfYgDeSkj)
705 [oYfYgDeSkj](https://openreview.net/forum?id=oYfYgDeSkj).
- 706 Johannes Welbl, Nelson F. Liu, and Matt Gardner. Crowdsourcing multiple choice science ques-
707 tions. In Leon Derczynski, Wei Xu, Alan Ritter, and Tim Baldwin (eds.), *Proceedings of*
708 *the 3rd Workshop on Noisy User-generated Text*, pp. 94–106, Copenhagen, Denmark, Septem-
709 ber 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-4413. URL
710 <https://aclanthology.org/W17-4413/>.
- 711
- 712 Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu
713 Wang. Know your limits: A survey of abstention in large language models. *Transactions of the*
714 *Association for Computational Linguistics*, 13:529–556, 2025. doi: 10.1162/tacl.a.00754. URL
715 <https://aclanthology.org/2025.tacl-1.26/>.
- 716 Zhiqiu Xia, Jinxuan Xu, Yuqian Zhang, and Hang Liu. A survey of uncertainty estimation
717 methods on large language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova,
718 and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Lin-*
719 *guistics: ACL 2025*, pp. 21381–21396, Vienna, Austria, July 2025. Association for Computa-
720 tional Linguistics. ISBN 979-8-89176-256-5. URL [https://aclanthology.org/2025.](https://aclanthology.org/2025.findings-acl.1101/)
721 [findings-acl.1101/](https://aclanthology.org/2025.findings-acl.1101/).
- 722 Hongshen Xu, Zichen Zhu, Situo Zhang, Da Ma, Shuai Fan, Lu Chen, and Kai Yu. Rejection
723 improves reliability: Training LLMs to refuse unknown questions using RL from knowledge
724 feedback. In *First Conference on Language Modeling, 2024a*. URL [https://openreview.](https://openreview.net/forum?id=lJMioZBoR8)
725 [net/forum?id=lJMioZBoR8](https://openreview.net/forum?id=lJMioZBoR8).
- 726
- 727 Tianyang Xu, Shujin Wu, Shizhe Diao, Xiaoze Liu, Xingyao Wang, Yangyi Chen, and Jing Gao.
728 SaySelf: Teaching LLMs to express confidence with self-reflective rationales. In Yaser Al-
729 Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Em-*
730 *pirical Methods in Natural Language Processing*, pp. 5985–5998, Miami, Florida, USA, Novem-
731 ber 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.343.
732 URL <https://aclanthology.org/2024.emnlp-main.343/>.
- 733 Boyang Xue, Fei Mi, Qi Zhu, Hongru Wang, Rui Wang, Sheng Wang, Erxin Yu, Xuming Hu, and
734 Kam-Fai Wong. UAlign: Leveraging uncertainty estimations for factuality alignment on large
735 language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher
736 Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational*
737 *Linguistics (Volume 1: Long Papers)*, pp. 6002–6024, Vienna, Austria, July 2025. Association
738 for Computational Linguistics. ISBN 979-8-89176-251-0. URL [https://aclanthology.](https://aclanthology.org/2025.acl-long.299/)
739 [org/2025.acl-long.299/](https://aclanthology.org/2025.acl-long.299/).
- 740 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
741 Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu,
742 Jianwei Zhang, Jianxin Yang, Jiaxin Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu,
743 Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin,
744 Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yi-Chao Zhang,
745 Yunyang Wan, Yuqi Liu, Zeyu Cui, Zhenru Zhang, Zihan Qiu, Shanghaoran Qian, and Zekun
746 Wang. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.
- 747 Junxiao Yang, Jinzhe Tu, Haoran Liu, Xiaoce Wang, Chujie Zheng, Zhexin Zhang, Shiyao Cui,
748 Caishun Chen, Tiantian He, Hongning Wang, et al. Barrel: Boundary-aware reasoning for factual
749 and reliable lrms. *arXiv preprint arXiv:2505.13529*, 2025.
- 750 Hanning Zhang, Shizhe Diao, Yong Lin, Yi Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng
751 Ji, and Tong Zhang. R-tuning: Instructing large language models to say ‘I don’t know’. In
752 Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of*
753 *the North American Chapter of the Association for Computational Linguistics: Human Lan-*
754 *guage Technologies (Volume 1: Long Papers)*, pp. 7113–7139, Mexico City, Mexico, June
755 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.394. URL
<https://aclanthology.org/2024.naacl-long.394/>.

756 Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao,
757 Yu Zhang, Yulong Chen, et al. Siren’s song in the ai ocean: a survey on hallucination in large
758 language models. *arXiv preprint arXiv:2309.01219*, 2023.

759
760 Hang Zheng, Hongshen Xu, Yuncong Liu, Shuai Fan, Lu Chen, Pascale Fung, and Kai Yu. En-
761 hancing LLM reliability via explicit knowledge boundary modeling. In *Second Conference on*
762 *Language Modeling*, 2025. URL <https://openreview.net/forum?id=WLgfeRhuA0>.

763 Runchuan Zhu, Xinke Jiang, Jiang Wu, Zhipeng Ma, Jiahe Song, Fengshuo Bai, Dahua Lin, Li-
764 jun Wu, and Conghui He. GRAIT: Gradient-driven refusal-aware instruction tuning for ef-
765 fective hallucination mitigation. In *Findings of the Association for Computational Linguistics:*
766 *NAACL 2025*, pp. 4006–4021, Albuquerque, New Mexico, April 2025. Association for Compu-
767 tational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findings-naacl.223. URL
768 <https://aclanthology.org/2025.findings-naacl.223/>.

769 Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen
770 Zhang, Xinwei Long, Ermo Hua, Biqing Qi, Youbang Sun, Zhiyuan Ma, Lifan Yuan, Ning Ding,
771 and Bowen Zhou. Ttrl: Test-time reinforcement learning, 2025. URL [https://arxiv.org/](https://arxiv.org/abs/2504.16084)
772 [abs/2504.16084](https://arxiv.org/abs/2504.16084).

773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810 A THE USE OF LARGE LANGUAGE MODELS

811
812 We use LLMs to refine the usage of terminology and sentences. We also use LLMs to adjust para-
813 graph structure in the main body of this paper. We also use LLMs to search for related works.
814

815 B EXPERIMENT DETAILS

816 B.1 DATASET

817 We test on the following publicly available datasets.
818

- 819 • **Pararel** (Elazar et al., 2021) is a dataset of factual knowledge and relations that are originally for
820 evaluating masked language models. We follow Zhang et al. (2024), changing it into a question-
821 answering format for generation evaluation. The validation set contains 5,584 QA pairs, which
822 we use for evaluation.
- 823 • **TriviaQA** (Joshi et al., 2017) is a large-scale reading comprehension dataset containing over
824 95,000 trivia-style question–answer pairs. The validation set contains 11,313 QA pairs, which
825 we use for evaluation.
- 826 • **Natural Questions (NQ)** (Kwiatkowski et al., 2019) are aggregated queries issued to the Google
827 search engine. All questions of NQ can be answered using the contents of the English Wikipedia.
828 The validation set contains 3,610 QA pairs, which we use for evaluation.
- 829 • **SciQ** (Welbl et al., 2017) is a crowdsourced dataset of 13,679 science questions. It covers topics
830 from biology, chemistry, physics, and earth science from real educational materials. The valida-
831 tion set contains 1,000 QA pairs, which we use for evaluation.
832
833

834 B.2 DETAILS OF BASELINES

835 The core idea of existing abstention methods is to fine-tune LLMs to answer the questions within
836 their knowledge scope and abstain from those outside it. It consists of two steps.
837

838 First, given a training dataset D , these methods partition it into two subsets, the answerable set D_0
839 and the unanswerable set D_1 . For questions in D_0 , the model is assumed to be confident about it,
840 and the ground-truth answer is set to be the model’s standard response. For questions in D_1 , the
841 ground-truth text, it is changed to an abstention expression: “I don’t know”. R-Tuning (Zhang et al.,
842 2024) prompts LLMs to answer questions from the training set, then classifies correctly answered
843 questions into D_0 and incorrectly answered ones into D_1 . R-Tuning-U (Zhang et al., 2024) and
844 SE-Tuning (Tjandra et al., 2024) use uncertainty to partition the training set. R-Tuning adopts
845 predictive entropy, which clusters lexical equivalent answers together. SE-Tuning adopts semantic
846 entropy, which clusters semantically equivalent answers together. Let C_1, \dots, C_m be the classes that
847 were clustered from sampled responses; the entropy is given by:
848

$$849 \text{Entropy}(\mathbf{x}) \approx - \sum_{i=1}^m p_{\theta}(C_i | \mathbf{x}) \log p_{\theta}(C_i | \mathbf{x}) \approx - \sum_{i=1}^m \left(\frac{|C_i|}{M} \right) \log \left(\frac{|C_i|}{M} \right), \quad (10)$$

851 where M is the total number of samples. Then, D_0 and D_1 is obtained by a threshold τ . Then,
852 LLMs are fine-tuned on D_0 and D_1 , learning to minimize the following objective:
853

$$854 \mathcal{L}_{CE}(p_{\theta}) = - \sum_{x \in D_0, D_1} \sum_{t=1}^{|\mathbf{y}^{(x)}|} \log p_{\theta}(y_t^{(x)} | \text{prompt}, \mathbf{x}, \mathbf{y}_{<t}^{(x)}), \quad (11)$$

855 For consistency, GRPO-SE is trained on the same processed dataset as SE-Tuning. Its objective
856 function is Equation 6, enabling a controlled comparison of reward granularity and abstention be-
857 havior.
858

860 B.3 TRAINING DETAILS

861 We use 15,000 samples from the Pararel dataset for training. Before training, we remove questions
862 with semantic entropy below 1. These low-entropy samples lack sufficient semantic diversity, which
863

can bias the binary confidence reward and lead the model to abstain too infrequently. Moreover, their simplicity diminishes the relative advantage signal in GRPO training, reducing the effectiveness of policy optimization. In the semantic clustering, we concatenate the question and the answer as the input of the NLI model. we run each experiment with three different random seeds and report the mean performance.

B.4 HYPERPARAMETERS

For F1SCORE on Llama3-8B-Instruct and Qwen2.5-7B-Instruct, we use the same hyperparameter setups, which are shown in Table 6.

Table 6: Hyperparameter setups for GRPO trainer.

Parameter	Value
<i>General Settings</i>	
bf16	true
use_vllm	true
vllm_device	auto
vllm_enforce_eager	true
vllm_gpu_memory_utilization	0.7
vllm_max_model_len	4608
do_eval	false
<i>Training Configuration</i>	
gradient_accumulation_steps	4
gradient_checkpointing	true
gradient_checkpointing_kwargs	use_reentrant: false
learning_rate	3.0e-06
lr_scheduler_type	cosine_with_min_lr
lr_scheduler_kwargs	min_lr_rate: 0.1
warmup_ratio	0.05
num_train_epochs	1
per_device_train_batch_size	10
per_device_eval_batch_size	10
threshold of abstention	5
<i>Generation Settings</i>	
max_prompt_length	256
max_completion_length	64
num_generations	10
temperature	1.0
seed	42, 32, 22
<i>Logging and Saving</i>	
log_completions	true
log_level	info
logging_first_step	true
logging_steps	1
logging_strategy	steps
save_strategy	steps
save_steps	50
report_to	wandb
<i>Reward Configuration</i>	
reward_funcs	confidence, format, accuracy (semantic)
reward_weights	1.0 2.0, 4.0

C RELIABILITY METRIC

C.1 ANALYSIS FOR RELIABILITY SCORE

Xu et al. (2024a) propose the Reliability Score (RS), which encourages LLMs to offer maximal assistance while minimizing hallucinated errors. RS is defined as a weighted sum of accuracy and truthful rate, which is used by current studies (Xu et al., 2024a; Kim et al., 2025; Yang et al., 2025): $RS(\alpha) = \alpha \times \text{Truth.} + (1 - \alpha) \times \text{Acc.}$, the weight α is answering rate.

The components are:

$$\begin{aligned} N &= N_1 + N_2 + N_3 + N_4 + N_5 \\ \text{Acc.} &= \frac{N_1}{N} \\ \text{Truth.} &= \frac{N_1 + N_3 + N_5}{N} \\ \alpha &= 1 - \frac{N_3 + N_5}{N} = \frac{N_1 + N_2 + N_4}{N} \end{aligned}$$

Substituting these into the RS formula:

$$\begin{aligned} RS &= \left(\frac{N_1 + N_2 + N_4}{N} \right) \left(\frac{N_1 + N_3 + N_5}{N} \right) + \left(1 - \frac{N_1 + N_2 + N_4}{N} \right) \left(\frac{N_1}{N} \right) \\ &= \left(\frac{N_1 + N_2 + N_4}{N} \right) \left(\frac{N_1 + N_3 + N_5}{N} \right) + \left(\frac{N_3 + N_5}{N} \right) \left(\frac{N_1}{N} \right) \\ &= \frac{(N_1 + N_2 + N_4)(N_1 + N_3 + N_5) + N_1(N_3 + N_5)}{N^2} \end{aligned}$$

The partial derivative of RS with respect to N_4 is:

$$\frac{\partial RS}{\partial N_4} = \frac{N_3 + N_5 - (N_1 + N_2 + N_4)}{N^2} \quad (12)$$

The sign of this derivative is variable and contingent on the specific distribution of counts among the categories. It is not guaranteed to be negative. For instance, consider a perfect model evaluated on a dataset composed solely of unanswerable questions ($N_1 = N_2 = N_3 = 0$, with an initial state of $N_4 = 0$ and $N_5 = N$). The derivative at this point is:

$$\left. \frac{\partial RS}{\partial N_4} \right|_{N_4=0} = \frac{N_5 - N_4}{N^2} = \frac{N - 0}{N^2} = \frac{1}{N} > 0 \quad (13)$$

A positive derivative demonstrates that RS **perversely incentivizes** a perfect model to start making mistakes. An initial increase in N_4 from zero leads to an *increase* in the RS score. This behavior is antithetical to the fundamental goal of a reliability metric.

C.2 ANALYSIS FOR THE PROPOSED RELIABILITY F1

The proposed reliability $F1_{rel}$ is defined as:

$$F1_{rel} = \frac{2}{\frac{1}{F1_{ans}} + \frac{1}{F1_{abs}}}$$

First, we calculate the sum of the reciprocals:

$$\begin{aligned} \frac{1}{F1_{ans}} + \frac{1}{F1_{abs}} &= \frac{2N_1 + 2N_2 + N_3 + N_4}{2N_1} + \frac{N_3 + N_4 + 2N_5}{2N_5} \\ &= \frac{N_5(2N_1 + 2N_2 + N_3 + N_4) + N_1(N_3 + N_4 + 2N_5)}{2N_1N_5} \\ &= \frac{2N_1N_5 + 2N_2N_5 + N_3N_5 + N_4N_5 + N_1N_3 + N_1N_4 + 2N_1N_5}{2N_1N_5} \\ &= \frac{4N_1N_5 + 2N_2N_5 + N_1N_3 + N_1N_4 + N_3N_5 + N_4N_5}{2N_1N_5} \end{aligned}$$

Finally, substitute this result back into the formula for $F1_{rel}$:

$$\begin{aligned} F1_{rel} &= \frac{2}{\frac{4N_1N_5 + 2N_2N_5 + N_1N_3 + N_1N_4 + N_3N_5 + N_4N_5}{2N_1N_5}} \\ &= \frac{4N_1N_5}{4N_1N_5 + 2N_2N_5 + N_1N_3 + N_1N_4 + N_3N_5 + N_4N_5} \end{aligned}$$

In stark contrast, the partial derivative of $F1_{rel}$ with respect to N_4 is:

$$\frac{\partial F1_{rel}}{\partial N_4} = -\frac{F1_{rel}^2}{4} \left(\frac{1}{N_1} + \frac{N_3 + N_4 + N_5}{N_5^2} \right) \quad (14)$$

The sign of this derivative is **consistently non-positive**. We can analyze its components to confirm this:

- The term $F1_{rel}^2$ is always non-negative.
- The term within the parentheses, $\left(\frac{1}{N_1} + \frac{N_3 + N_4 + N_5}{N_5^2} \right)$, is a sum of non-negative quantities and is strictly positive for any non-trivial case where the metric is defined ($N_1 > 0$ and $N_5 > 0$).
- The leading factor of $-\frac{1}{4}$ ensures the entire expression is non-positive.

This non-positive derivative signifies the ideal behavior. It mathematically guarantees that any increase in the critical error count N_4 will **always lead to a decrease or no change** in the $F1_{rel}$ score. The metric consistently and correctly penalizes unreliability.

Conclusion. The derivative analysis provides a rigorous mathematical basis for concluding that $F1_{rel}$ is a superior metric for evaluating model reliability. While the RS metric can fail catastrophically by rewarding incorrect behavior, the $F1_{rel}$ metric provides a robust, reliable, and interpretable gradient that always punishes critical errors. This makes $F1_{rel}$ a far more suitable and trustworthy measure for developing reliable AI systems.

C.3 SPECIAL CASE

To highlight the superiority of our proposed metric $F1_{rel}$ (the harmonic mean of $F1_{ans}$ and $F1_{abs}$) over the Reliability Score (RS), we analyze a specific, illustrative scenario: a dataset composed exclusively of unanswerable questions. A robust reliability metric must correctly identify and reward a model that abstains from such questions, while penalizing a model that makes incorrect guesses. Consider a dataset that contains only unanswerable questions. This implies:

$$N_1 = 0, \quad N_2 = 0, \quad N_3 = 0$$

the total number of samples is therefore $N = N_4 + N_5$.

We evaluate two distinct models on this dataset:

- **Model A (The Perfect Abstainer):** This model performs perfectly by correctly identifying all questions as unanswerable and abstaining every time, defined by $N_5 = N, N_4 = 0$.
- **Model B (The Reckless Guesser):** This model is flawed. It abstains half the time but provides an incorrect answer for the other half, defined by $N_5 = 0.5N, N_4 = 0.5N$.

Our expectation: A valid metric should award Model A a perfect (or maximum) score and Model B a significantly lower score.

Analysis of the Reliability Score (RS) The simplified formula for RS in this case is:

$$RS = (\alpha \cdot \text{Cov}) + ((1 - \alpha) \cdot \text{Acc}) = \left(\frac{N_4}{N} \cdot \frac{N_5}{N} \right) + \left(\left(1 - \frac{N_4}{N}\right) \cdot 0 \right) = \frac{N_4N_5}{N^2}.$$

For Model A:

$$RS_A = \frac{(0)(N)}{N^2} = 0$$

1026 For Model B:

1027

1028

1029

1030

$$RS_B = \frac{(0.5N)(0.5N)}{N^2} = \frac{0.25N^2}{N^2} = 0.25$$

1031

RS erroneously assigns Model A a score of 0 and Model B a higher score of 0.25.

1032

1033

Evaluation of the proposed $F1_{rel}$ Metric

1034

In this scenario, $F1_{rel}$ depends solely on $F1_{abs}$, reflecting abstention performance:

1035

1036

1037

1038

$$F1_{rel} = F1_{abs} = \frac{2N_5}{N_3 + N_4 + 2N_5}$$

1039

For Model A:

1040

1041

1042

1043

1044

$$F1_{absA} = \frac{2(N)}{0 + 0 + 2(N)} = \frac{2N}{2N} = 1$$

1045

This assigns a perfect score of 1, reflecting optimal performance. For Model B:

1046

1047

1048

1049

$$F1_{absB} = \frac{2(0.5N)}{0 + 0.5N + 2(0.5N)} = \frac{N}{0.5N + N} = \frac{N}{1.5N} = \frac{2}{3} \approx 0.67$$

1050

This appropriately penalizes Model B for incorrect answers.

1051

1052

Conclusion

1053

The Reliability Score (RS) exhibits a critical flaw by rewarding incorrect answers on unanswerable questions, contradicting the principles of reliability. In contrast, $F1_{rel}$, via its $F1_{abs}$ component, accurately assigns Model A a perfect score and penalizes Model B, demonstrating its robustness and suitability as a reliability metric.

1057

1058

1059

D EVALUATION PROMPT

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

ICL Prompt

Answer the following questions as briefly as possible.

Question: {demo question 1}

Answer: {demo answer 1}

Question: {demo question 2}

Answer: {demo answer 2}

...

Question: {input question}

Answer:

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

ICL-IDK Prompt

Answer the following questions as briefly as possible. If you don't know the answer, please simply say "I don't know."

Question: {demo question 1}
Answer: {demo answer 1}

Question: {demo question 2}
Answer: I don't know.

...

Question: {input question}
Answer:

ICL-Unsure Prompt

You are a helpful and truthful AI Assistant. You should answer the question as briefly as possible, then present your confidence. If you are sure about your answer, please say "I am sure" after your answer; otherwise, say "I am unsure".

Question: {demo question 1}
Answer: {demo answer 1} I am sure.

Question: {demo question 2}
Answer: {demo answer 2} I am unsure.

...

Question: {input_question}
Answer:

R-Tuning, R-Tuning-U and SE-Tuning Prompt

Answer the following questions as briefly as possible. If you don't know the answer, please simply say "I don't know."

Question: {demo question 1}
Answer: {demo answer 1}

Question: {demo question 2}
Answer: I don't know.

...

Question: {input question}
Answer:

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

```

GRPO-SE and FiSCORE Prompt

You are a helpful and truthful AI Assistant that provides responses include
answer and confidence. You first answer the question as briefly as possible
enclosed by <answer> and </answer> and then provide your confi-
dence in sure or unsure about the answer enclosed by <confidence>
and </confidence>. Respond in the following format:
<answer>...</answer>
<confidence> sure or unsure </confidence>

Question: {demo question 1}
Answer: <answer> {demo answer 1} </answer>
<confidence> sure </confidence>

Question: {demo question 2}
Answer: <answer> {demo answer 2} </answer>
<confidence> unsure </confidence>

...

Question: {input question}
Answer:

```

E CASE STUDY

We present detailed examples using Llama3-8B-Instruct in Table 7. In question 1, both ICL-IDK and GRPO-SE exhibit over-abstention. While ICL achieves an accuracy of 70%, ICL-IDK generates two correct answers, one hallucinated response, and seven abstentions. GRPO-SE produces six correct answers but discards five of them. In contrast, FiSCORE successfully retains all correct answers while discarding all incorrect ones. In question 2, GRPO-SE demonstrates overaggressive behavior by retaining all incorrect answers, whereas FiSCORE appropriately discards three incorrect responses, achieving calibrated post-hoc abstention.

F ABSTENTION RATE

Table 8 shows the abstention rate for each method on each dataset.

G COMPUTATION COST

R-Tuning (Zhang et al., 2024) trains LLMs to respond with “I don’t know” explicitly when feeded a uncertain question. In contrast, FiSCORE enables self-reflective abstention or answering, at the cost of increased computation. We compare ICL, R-Tuning, and FiSCORE in terms of the average generation length of all questions in the test sets and the generation time on 100 randomly sampled questions using an NVIDIA L40 GPU. As shown in Table 9, R-Tuning exhibits slightly slower generation speed than ICL, though in certain cases it performs comparably or even faster. In contrast, FiSCORE incurs a 3-4x increase in generation time relative to ICL. This overhead arises from the structured output format: for answered questions, our method introduces additional tokens for tags and confidence expression; for abstained questions, it further generates a full-length answer to justify the refusal, thereby increasing computational cost.

H LIMITATIONS

Our method has some limitations. First, we must manually set the threshold τ during GRPO training. Second, the $F1_{ans}$, $F1_{abs}$, and $F1_{rel}$ metrics are only suitable for evaluating alignment methods,

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Table 7: Case studies of different methods on Llama3-8B-Instruct. We sample 10 answers for each question with the temperature of 0.2.

Question1: The world’s largest marketer of fruit juices, what is the juice arm of the Coca Cola company? Answer: maid			
ICL	ICL-IDK	GRPO-SE	FISCoRE
minute maid	maid	<answer> maid </answer> <confidence> sure </confidence>	<answer> maid </answer> <confidence> sure </confidence>
minute maid	maid	<answer> maid </answer> <confidence> unsure </confidence>	<answer> maid </answer> <confidence> sure </confidence>
minute maid	simply	<answer> maid </answer> <confidence> unsure </confidence>	<answer> maid </answer> <confidence> sure </confidence>
minute maid	I don’t know.	<answer> maid </answer> <confidence> unsure </confidence>	<answer> maid </answer> <confidence> sure </confidence>
minute maid	I don’t know.	<answer> maid </answer> <confidence> unsure </confidence>	<answer> maid </answer> <confidence> sure </confidence>
minute maid	I don’t know.	<answer> maid </answer> <confidence> unsure </confidence>	<answer> maid </answer> <confidence> sure </confidence>
minute maid	I don’t know.	<answer> simply </answer> <confidence> unsure </confidence>	<answer> simply </answer> <confidence> unsure </confidence>
simply	I don’t know.	<answer> simply </answer> <confidence> unsure </confidence>	<answer> simply </answer> <confidence> unsure </confidence>
fanta	I don’t know.	<answer> fanta </answer> <confidence> unsure </confidence>	<answer> fanta </answer> <confidence> unsure </confidence>
fanta	I don’t know.	<answer> fanta </answer> <confidence> unsure </confidence>	<answer> fanta </answer> <confidence> unsure </confidence>
Question2: Which mythical beings were said to inhabit the slopes of Mount Etna? Answer: Cyclops			
ICL	ICL-IDK	GRPO-SE	FISCoRE
Cyclops	Cyclops	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Cyclops	Cyclops	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Cyclops	Cyclops	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Cyclops	Cyclops	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Cyclops	Cyclops	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Cyclops	Cyclops	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Cyclops	I don’t know.	<answer> Cyclops </answer> <confidence> sure </confidence>	<answer> Cyclops </answer> <confidence> sure </confidence>
Centaurs	I don’t know.	<answer> Centaurs </answer> <confidence> sure </confidence>	<answer> Centaurs </answer> <confidence> unsure </confidence>
Centaurs	I don’t know.	<answer> Centaurs </answer> <confidence> sure </confidence>	<answer> Centaurs </answer> <confidence> unsure </confidence>
Centaurs	Centaurs	<answer> Centaurs </answer> <confidence> sure </confidence>	<answer> Centaurs </answer> <confidence> unsure </confidence>

Table 8: Abstention rate

	Llama3-8B-Instruct				Qwen-2.5-7B-Instruct			
	Pararel	TriviaQA	NQ	SciQ	Pararel	TriviaQA	NQ	SciQ
ICL-IDK	48.30	12.35	55.79	15.40	54.91	21.97	28.67	4.10
ICL-Unsure	37.23	17.86	7.26	4.60	56.45	18.99	21.66	11.60
R-Tuning	54.03	59.69	70.69	75.30	66.47	84.10	92.85	78.00
R-Tuning-U	56.86	73.62	75.32	55.50	77.69	85.38	90.01	76.70
SE-Tuning	43.81	63.18	71.16	83.90	69.03	88.78	88.34	58.51
GRPO-SE	55.93	35.76	68.59	24.20	63.53	40.08	62.91	22.30
FiSCORE	21.25	23.03	58.25	28.50	28.28	35.79	59.97	24.90

Table 9: The results of the computation cost. Left: average number of tokens of the model’s response. Right: average generation time of 100 random questions.

Model	Avg. tokens				Avg. inference time (s)			
	Pararel	TriviaQA	NQ	SciQ	Pararel	TriviaQA	NQ	SciQ
Llama3-8b-Instruct								
ICL	2.86	3.93	5.14	3.47	7.75	11.93	12.36	10.39
R-Tuning	5.33	5.74	6.47	3.97	13.84	16.14	15.31	20.77
FiSCORE	15.83	16.82	17.70	16.06	31.79	35.29	35.13	32.82
Qwen2.5-7B-Instruct								
ICL	3.36	5.54	12.44	3.64	12.23	12.49	19.49	11.04
R-Tuning	4.35	4.92	5.84	3.41	13.44	15.54	16.62	20.73
FiSCORE	12.17	16.84	19.36	15.29	42.69	44.54	54.20	42.49

which means that we must know the knowledge that the model possesses before alignment. Future evaluation metrics for LLMs should encourage models to acknowledge their knowledge gaps rather than guessing answers. Our method teaches LLMs to first generate an answer and then express the abstention decision, which requires more computation than directly abstaining from answering.